JOURNAL OF APPLIED BEHAVIOR ANALYSIS 1994, 27, 739–760 NUMBER 4 (WINTER 1994)

# *SHAPING IN THE 21ST CENTURY: MOVING PERCENTILE SCHEDULES INTO APPLIED SETTINGS*

GREGORY GALBICKA

WALTER REED ARMY INSTITUTE OF RESEARCH

The present paper provides a primer on percentile reinforcement schedules, which have been used for two decades to study response differentiation and shaping in the laboratory. Arranged in applied settings, percentile procedures could be used to specify response criteria, standardizing treatment across subjects, trainers, and times to provide a more consistent training environment while maintaining the sensitivity to the individual's repertoire that is the hallmark of shaping. Percentile schedules are also valuable tools in analyzing the variables of which responding is a function, both inside and outside the laboratory. Finally, by formalizing the rules of shaping, percentile schedules provide a useful heuristic of the processes involved in shaping behavior, even for those situations that may not easily permit their implementation. As such, they may help further sensitize trainers and researchers alike to variables of critical importance in behavior change.

DESCRIPTORS: shaping, response differentiation, percentile schedules, reinforcement density, operant conditioning

In behavior analysis, it is often desirable to take a behavioral repertoire and mold it into something different. In developmental disabilities, self-care, social, and vocational skills often need to be trained; in sports psychology, more skilled performance is a frequent goal; and education itself is nothing but the modification of behavioral repertoires. When dealing with operant behavior, this change is generally effected via a process termed *shaping,* a shorthand for differential reinforcement of successive approximations to a terminal response (see Skinner, 1953). Organisms and environments continuously shape the behavior of other organisms by providing consequences differentially following particular responses demonstrating certain criterion characteristics. A response (e.g., adding some wine to the spaghetti sauce) followed by a positive reinforcer (e.g., a better tasting sauce that wins the approval of your dinner companions) will increase in frequency over one provided no consequences, changing the response distribution to include relatively more responses similar to the type reinforced. Extinction, on the other hand, not only decreases response frequency but also temporarily increases the variability in responding, thereby increasing the probability that a response from the reinforced class will occur. Shaping occurs when reinforcement and extinction are used in combination with a systematically changing set of response criteria to reinforce responding differentially (i.e., to reinforce responses exhibiting some criterional attribute while not reinforcing noncriterional responses). Thus, the successful shaper must carefully ascertain characteristics of an individual's present response repertoire, explicitly define characteristics the final behavior will have at the end of training, and plot a course between reinforcement and extinction that will bring the right responses along at the right time, fostering the final behavioral sequence while never losing responding altogether.

For such a prevalent technique, shaping is subject to considerable variation between subjects, between trainers, and even with a single subject–trainer pair at different times. The "rules" of shaping are typically qualitative in nature only, with little empirical data on the effects of quantitative variation. As such, the rules constitute more an art form than a science, and the attitude is often that shaping is something you can only learn by doing—it is con-

I thank Timothy F. Elsmore, Carol Pilgrim, G. Jean Kant, and members of the Physiology and Behavior Branch for comments on an earlier version of this report. The views of the authors do not purport to reflect the position of the Department of the Army or the Department of Defense (para 4-3, AR 360-5).

Request reprints from the author, Department of Medical Neurosciences, WRAIR, Washington, D.C. 20307-5100.

tingency shaped, not rule governed. The contingencies that shape effective shaping are themselves found in the effectiveness of interactions between trainer and client, and will necessarily vary with a change in either or both of the individuals. Hence, developing a quantitative science of shaping may seem, if not beyond reach, at least difficult to the point of having little applied relevance.

The present paper argues against this view. It presents a primer on the workings of percentile reinforcement schedules, procedures that have been used in laboratory studies of response differentiation and shaping for over two decades. Percentile schedules disassemble the process of shaping into its constituent components, translate those components into simple, mathematical statements, and then use these equations, with parameters specified by the experimenter or trainer, to determine what presently constitutes a criterional response and should therefore be reinforced.

Percentile schedules, however, do more than automate shaping. In addition, they make explicit and objective the criteria that define responses as criterional or noncriterional throughout acquisition and maintenance, providing explicit prior control over reinforcement density as well as criterional response probability. Because of this, they provide almost complete independence from trainer- and subject-related variables. This allows all subjects to be trained in a specified manner despite changes in the trainer or the subject, or at different points in the differentiation.

### *Shaping's Golden Rules*

As a prelude to presenting the mechanics of percentile schedules, it may be helpful to consider the verbal "rules" of shaping as they have been distilled from experience, and as they are generally presented to students and trainers when teaching the fundamentals of shaping. This presentation is explicitly rudimentary, in that it attempts to provide a common starting point for the discussion of how to transform the verbal rules into more explicit, quantitative ones. A review of these rules clarifies the operation of percentile schedules by providing a point of correspondence between what we already know verbally about shaping and what percentile schedules provide in the form of equations.

Shaping involves differential reinforcement of operant behavior. Because a behavior must occur prior to being reinforced, the first rule of shaping is generally some variant of "Start where the subject is"—set the initial reinforcement criterion at a value within the subject's current repertoire. That is, the current repertoire will be characterized by a distribution of responses varying across some range of values. Requiring values completely outside this distribution at the beginning of training transforms shaping into extinction, because all responses emitted will fail to meet the reinforcement criterion and thus will not produce reinforcement.

The next rule is generally of the form "Clearly define the terminal response." This ensures that we know when the differentiation has been successful and also often helps to define important behavioral dimensions and potential intermediate steps. Rule 1 provides a clear understanding of the subject's initial behavioral distribution. Rule 2 specifies characteristics of the ultimate distribution once shaping is complete. By comparing attributes of the initial and terminal responses for the ways in which they differ, an idea of the kinds of response characteristics that should be measured will emerge. For example, suppose I am interested in becoming a long-distance runner. Given my previous interest in running (zero), it seems to be a good idea to establish explicit reinforcement contingencies external to the joy of running per se to shape running. Setting aside questions of finding a suitable reinforcer and someone to deliver them appropriately, the main problem is to develop a program of differential reinforcement contingencies that will result in my ultimately emitting the terminal response of completing a marathon. I can probably dispense immediately with measuring sit-ups and concentrate on running, because the dimension of interest has something to do with running. I also need not concern myself with running speed, only stamina (i.e., I am not foolish enough to want to *win* my first marathon, only to finish it). Further, focusing on the terminal response helps to define the response unit as "running 26 miles and 385 yards," emphasizing that distance is the functional aspect of a running episode being differentiated, not other possible aspects or units (e.g., number of strides). Other units could be shaped without necessarily achieving the terminal response of completing a marathon (e.g., number of strides can increase with little change in distance if strides shorten). Emphasizing the difference in distance between the initial and terminal runs forces that dimension to be the functional unit of behavior. Finally, by noting the difference between the current level and the goal, a number of finite criteria can be defined that increase the required distance run in fixed arithmetic or exponential increments (i.e., adding or multiplying a constant to each level to generate a series of intermediate values). All this follows as a consequence of specifying in advance where responding is (Rule 1) and where it will end (Rule 2).

The third rule is "Use small steps." To use the running example again, this rule indicates that the smaller the increment in the reinforcement criterion (i.e., distance run) at each criterion change, the less likely it will be that responding will reach a point at which the variation from instance to instance will not include enough reinforceable values to maintain a fair degree of behavior. Compare two training regimens, one of which increases the reinforcement criterion in increments of half a mile and the other of which increases it by 5 miles at each criterion shift. With each change in the criterion, the natural variation in running stamina all but guarantees that the new criterion will be met by a run in the near future under the former regimen, whereas requiring a more substantial change (the latter regimen) decreases the probability that the criterion will be met following each change.

The last rule, not always taught explicitly, is "Reinforce movement, not position." Criteria established in terms of the *change* they generate will more likely result in behavior change than those anchored to some static quality or product. For the running example, this rule suggests that reinforcement contingent on a specified *increase in distance* from the previous run (i.e., $x$ miles *farther*) will be more successful in increasing running than are criteria set to specific distances. Criteria that emphasize behavior change may increase the probability that behavior will change when the criterion is again shifted, increasing the likelihood that subsequent criteria will also be met.

### *Shaping's Foundations*

These rules work because of the manner in which reinforcement and extinction, the component processes of differential reinforcement, work. In Skinner's (1938) earliest writings on operant behavior, he noted that reinforcement increased the rate of a class of responses, not the rate of a particular response. Members of this class vary with respect to the exact distribution of any of a number of measurable response characteristics (e.g., location, intensity, duration, topography, etc.), but are invariant with respect to their function—they all produce the consequence in question. If that consequence reinforces the operant, similar responses will more likely recur.

Extinction is often presented as the opposite of reinforcement, but it is very much more. No longer reinforcing an operant ultimately does decrease the rate of the response class. If this was the only effect of extinction, however, learning would be very constrained. Because reinforcement typically generates responses similar to those previously reinforced, some other mechanism must generate novel behaviors in response to a changing criterion. That mechanism is extinction. Removing reinforcement initially generates variability in behavior (see Galbicka, 1988, for a review of the experimental literature). As the patterns previously learned begin to extinguish, they recombine with other response units occasioned by the same environment; oftentimes previously trained units reemerge (e.g., Epstein, 1983) or other, previously ineffective, sources of control generate novel responses. This variability is important because it increases the probability that a response meeting the new criteria will be emitted. When we check into the hotel for a convention and are confronted by a bathroom fixture we have never seen before, we first behave in a fashion appropriate to the one at home. Turn the handle, or twist the dial—but if none of those work, we slowly start trying other responses. Where do they come from? They are

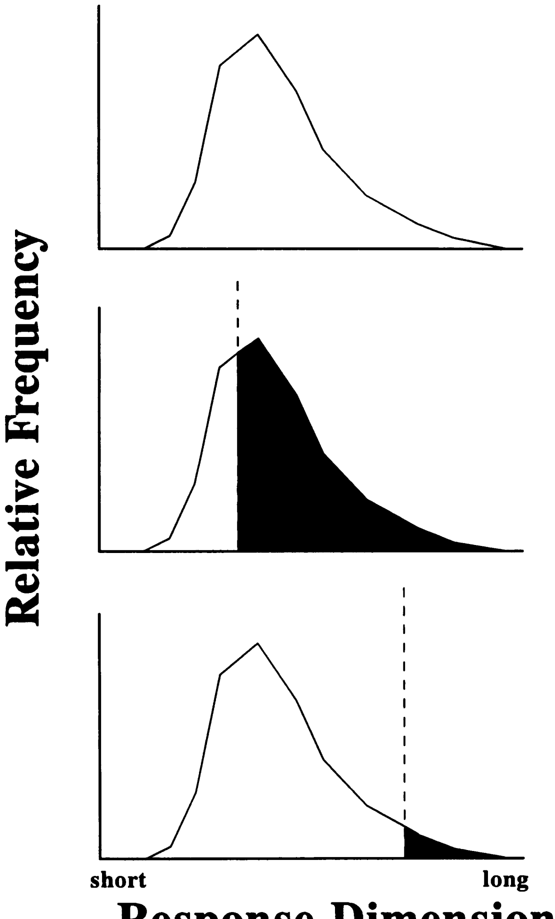


Figure 1. Illustration of how different criteria (the vertical lines in the bottom two panels) applied to a single response distribution establish different overall reinforcement densities and differences in the degree to which reinforcement is correlated with extreme values.

extinction-induced variations in responding that relate to our past histories with respect to buttons, levers, dials, and so forth. Push it, punch it, turn it, just keep varying—sooner or later something will work. Reinforcement, then, generates responses that are identical in function and similar in appearance to those preceding it. Extinction, for a while at least, is an aid to learning, because it generates a local high rate of variable behavior that can come under control of the changing contingencies that define a differentiation. However, if this transient increase in variability does not successfully induce a member of the new criterion class, extinction ultimately will eliminate responding.

The temporary nature of this effect is the conceptual basis of Rule 3. Step size determines the probability of a criterional response, which in turn determines when reinforcers are presented. Suppose that the top distribution in Figure 1 refers to the current distribution of a behavior, and suppose that we wish to generate longer values. (The actual response as well as the values represented along the response dimension are irrelevant for now; we simply want to increase the frequency of "longer" values and thereby shift the entire distribution to the right.) As a first step towards this end, we could decide to impose the criteria indicated by the vertical line in the middle panel and then reinforce only responses longer than that value (i.e., the shaded portion). Alternatively, the criterion could be set at the value indicated by the dashed line in the lower panel. The criterion in the middle panel is relatively lax, in that a majority of responses observed during the baseline would exceed the criterion. If behavior does not change (i.e., if the distribution of behavior remains constant), most responses will still produce reinforcement (the shaded portion of the distribution). The lower panel shows a much more stringent criterion; much longer values are required for reinforcement, and if the distribution remains constant only a small proportion of responses are reinforced. Which step size will lead to the most rapid shaping? There are advantages and disadvantages of either selection. The criterion in the middle panel protects against the complete elimination of behavior by all but guaranteeing a relatively high reinforcement density after the criterion is put into effect. However, that protection comes at the expense of differential reinforcement—the range of values reinforced is very large and includes many relatively short and medium values as well as long ones. As such, it is not likely that very long values will soon begin to predominate. The advantage of highly differential reinforcement resides with the stricter criterion depicted in the lower panel. Because only relatively large-valued responses produce reinforcement, similar large-valued ones will more likely recur once reinforced. The disadvantage is that responding could extinguish altogether before a criterional response occurs.

Deciding between these two alternatives is the crux of shaping; establishing criteria that provide sufficient but not excessive reinforcement is central to the success of the procedure. Less often appreciated is the importance of shifting criteria at the right time. Assume for the moment that we somehow solve the dilemma and set the criterion at the point denoted by the dashed vertical line in the top panel of Figure 2. If all baseline responses were reinforced, then imposing the criterion will result in an immediate substantial decrease in the density of reinforcement (i.e., only the initially small proportion of responses above the criterion will produce reinforcement). This partial extinction is important in producing the local effects noted above, ultimately generating greater values that will exceed the criterion and be reinforced. After exposure to the criterion for some period of time, responding might come to resemble the distribution depicted in the middle panel. After even further exposure, it might resemble that shown in the lower panel. The criterion (i.e., the vertical line) remains unchanged in each panel, but, because of the progressive change in the distribution, more responses meet criterion with extended exposure (approximately one half of the distribution meets criterion in the middle panel, whereas practically all responses exceed the criterion in the lower panel). When should the next level be imposed? That is, when should the criterion be shifted towards even greater values?

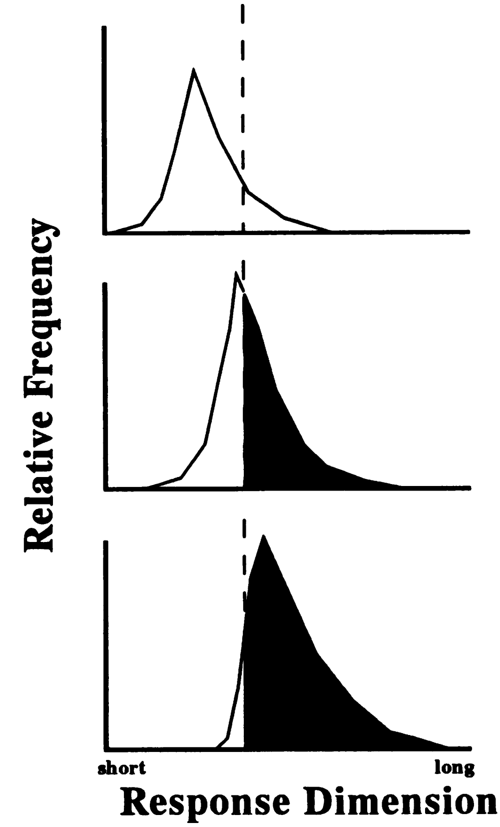


Figure 2. Three hypothetical distributions illustrating responding under baseline (top) and after exposure to a particular differential reinforcement criterion (i.e., responses longer than the value indicated by the vertical line produce reinforcement) for a moderate period (middle panel) or for a more extended period (lower panel).

The conservative approach might be to provide the client with substantial training and ensure a high probability of long responses before increasing the criterion. However, consider what happens to reinforcement frequency each time the criterion is changed. Because the criterion by definition includes only a portion of the current distribution, imposing a new criterion is always associated with a decrease in reinforcement density. Only by shifting the distribution of responses emitted to even longer values can the reinforcement density be returned to its former level. A plot of reinforcement density across time would reveal a pattern like a sawtooth; with each change in the criterion, reinforcement density drops abruptly, but as behavior gradually changes to include more and more criterional responses, reinforcement density gradually increases until the cycle repeats with the next criterion change. This cyclic change in reinforcement density is more pronounced following extended training. That is, reinforcement density can fall from a maximum probability of about .5 if the criterion is changed when the distribution of responding resembles that shown in the middle panel of Figure 2. Extended training provides a higher reinforcement probability (almost 1.0; see the lower panel of Figure 2) prior to the criterion shift, and hence a greater potential reduction in reinforcement frequency from that value once the criterion changes.

This discussion suggests a corollary to Rules 1 and 3, namely that criteria should be adjusted often to remain sensitive to current behavior and keep reinforcement differential and intermittent. Consistently intermittent reinforcement is essential to successful differentiation because it generates persistence in the face of extinction. The experimental and applied literatures are both replete with demonstrations that a prior history of intermittent reinforcement generates far more responding during extinction than does a history of continuous reinforcement. Hence, rather than wait for most responses to meet criterion and then drastically reducing reinforcement frequency by shifting criteria infrequently, it is better to change criteria frequently to maintain both a relatively *constant* reinforcement density and an *intermittent* one. Both characteristics decrease the likelihood of losing control over responding prior to the acquisition of the terminal response.

### *Percentile Schedules: Formalized Shaping*

The preceding discussion suggests that any attempt to formalize these rules into a procedure should include the following characteristics: (a) It should set criteria relative to current behavior and change them rapidly as behavior changes. (b) It should establish criteria in such a way that some sufficiently large proportion of responses is reinforced, but that proportion cannot be so large as to dilute the differential nature of the contingency. (c) It should provide reinforcement consistently and intermittently, despite any changes in behavior upon which that reinforcement ultimately depends. (d) Finally, it should provide some terminal response definition.

The third characteristic is the most problematic in formalizing shaping; the procedure must provide a consistent, intermittent density of reinforcement. The traditional view of shaping holds that responding is a *dependent* variable subject to change and not an independent one that can be controlled prior to its occurrence. Yet, reinforcing only responses in the criterion zone (required by the second characteristic above) while keeping reinforcement density constant at some specified intermittent value seemingly requires prior knowledge of, and control over, the proportion of criterional responses. For example, if the desired probability of reinforcement is .25, and all criterional responses are reinforced, then the criterion must be set so that 25% of all responses exceed the criterion. This is contrary to the role normally attributed to responding in a shaping procedure, where changes in the probability of criterional responses define the effectiveness of a differentiation rather than being subject to experimental control.

How can responding be characterized so that the probability of criterional and noncriterional responses could be specified in advance? The percentile solution, developed and expanded by Platt (1973) and colleagues, is momentarily to abandon the exact physical characteristics of the response and treat it as an ordinal quantity. Ordinal quantities are values that carry only an associated rank, as opposed to the more typical means of quantifying observations by assigning a cardinal number and a standard unit. For example, height can be cardinally measured as 6 ft 8 in., or it can be ordinally measured by comparing two people and placing the taller one on the left.

We generally prefer to use the cardinal method because it provides more detailed information. However, there are times when measuring things ordinally has its advantages. Suppose the horizontal line in the top panel of Figure 3 represents a dimension along which a particular response of interest may vary. Using the running example again, the response is an episode of running and the dimension is the distance run. The probability that the next response (a run) will fall somewhere along this dimension is 1 (all runs have some distance). Suppose that the first run sampled falls at Point A, depicted in the second panel of the figure. How likely is it that the next run will be less than A, and how likely is it to exceed A? Intuition may suggest that the next run is equally likely to fall above or below the single observation at A, and that is in fact the case. Suppose the next run has the value represented by the point labeled B in the third line. With what probability will the run after that fall into each of the three intervals bounded by the two observations (i.e., below A, between A and B, and above B)? Now intuition suggests that values less than A or greater than B should be observed proportionately more often than ones between A and B, because the interval AB is smaller than the other two intervals. In fact, a subsequent response will again fall into each interval with equal probability, or one third of the time. Adding a third observation, at C in the fourth panel, generates four intervals by splitting the interval from A to B into the intervals AC and CB. Although these two intervals are clearly not the same size, the probability of the next observation falling into each interval bounded by C, and the other two intervals as well, is one fourth.

The generalization derived from the above example is that $m$ previous observations create $m + 1$ intervals, one of which must contain the next observation. This fact alone is insufficient to derive percentile procedures. A single, simple constraint must be attached: Observations must be sampled randomly and independently from the population of values. This means that knowing the value of the current response cannot help to predict the occurrence of a later one. If these two conditions are met, the probability of a criterional response can actually be predicted and controlled (i.e., specified in advance). Even when this assumption is clearly violated, some steps discussed below can be taken to maintain control over reinforcement probability.

The counterintuitive notion that intervals of different sizes are equally likely to contain the next observation arises because the line represents a *cardinal* scale, but the question of which interval will contain the next observation relates to the *ordinal* properties of the observations. For the moment, ignore the fact that there are physical values attached to any of these observations, and treat them solely in terms of their ranks. In any distribution of values, there is one and only one value ranked 1st, 2nd, 3rd, and so forth. The question of interest is not "What is the expected value of the next observation (i.e., what distance will next be run)?" but rather is "Where will the next observation rank?" If the assumption of independence is met, it will be as likely to rank first or last or anywhere in between, depending on the number of prior observations. The probability that it will rank between 1 and $m + 1$ is 1 (i.e., it must fall somewhere on the line). This also equals the sum of the probabilities that it falls into each of the intervals defined by each observation. That is, the bottom line is segmented into four parts by the three prior observations A, B, and C. If the next observation fell below A, it would receive a rank of 1st (lowest); if it fell between A and C it would net a rank of 2nd lowest, between C and B a rank equal to 3rd lowest, and above B a rank of 4th lowest. Each of these rankings is equally likely by definition. Further, the sum of the probabilities of ranking 1st, 2nd, 3rd, or 4th must equal 1 (i.e., no other ranks exist). Hence, each rank will occur with a probability equal to the reciprocal of the number of intervals available to contain the next response. Thus, given any distribution of $m$ prior observations, the probability the next observation will fall into each one of the $m + 1$ intervals bounded by

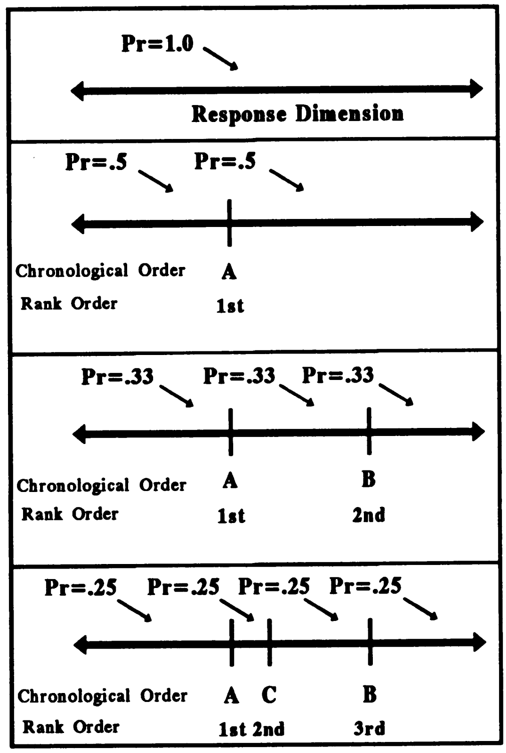


Figure 3. Illustration of the how successive observations (denoted A, B, etc.) divide a dimension into intervals having equal probabilities of containing the next response. See text for details.

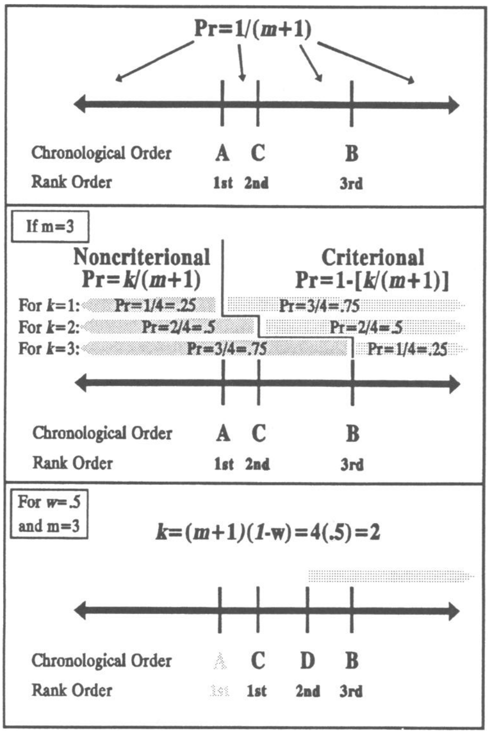


Figure 4. Upper panel: generalization of the effect illustrated in Figure 3. Middle panel: illustration of how criteria can be established in terms of exceeding particular ranks, and the resulting effects on the probability of observing a criterional response. Lower panel: inversion of the middle panel, demonstrating the percentile equation, which determines the rank ($k$) that must be exceeded in the current distribution of $m$ responses to observe a criterional response with the specified probability $w$.

these observations is $1/(m + 1)$, as shown for $m = 1$, 2, or 3 in the bottom three panels of Figure 3, respectively, or for the generalized case in the top panel of Figure 4.

The relevance to shaping depends on understanding one implication of the above. If the probability that the next observation will be ranked in any one interval defined by $m$ observations is $1/(m + 1)$, the probability that it will fall into any two of those intervals is twice that, or $2/(m + 1)$, the probability that it falls into any three would be $3/(m + 1)$, and so on. Hence, the probability that the next observation will fall into any one of $k$ intervals defined by $m$ observations is $k$ times the probability of falling into each interval, or $k/(m + 1)$. This extension suggests the step taken in the second panel of Figure 4—establishing a criterion at the $k$th rank. That is, rather than setting the criterion at a particular fixed, physical value, the criterion can specify that the next observation, to meet criterion, must rank higher than the value currently ranked $k$. The middle panel of Figure 4 illustrates the effects of establishing criteria at different values of $k$. When $k = 1$, responses will be considered criterional if they exceed the response currently ranked 1st (A). The probability of not meeting criterion equals the probability of falling into the one interval below A, which is $k/(m + 1) = 1/(3 + 1) = .25$. The probability of a criterional response (denoted $w$) is the complement, or $w = 1 - [k/(m + 1)] = .75$. Setting the criterion ($k$) so that the 2nd rank (C) must be exceeded raises the probability of not meeting criterion to the sum of the two intervals below C, or $2/(m + 1) = 2/4 = .5$. The probability of meeting it is again the complement, or $w = 1 - .5 = .5$. If all observations must be exceeded (i.e., $k = 3$), then the probability of a criterional response will be $w = 1 - [k/(m + 1)] = 1 - 3/4 = .25$. Thus, as the criterion is made more stringent (i.e., as $k$ is increased), the probability of observing a criterional response decreases accordingly, as intuition would suggest.

Notice that the rank and the chronological order of the observations are independent. In the sample shown in the figure, the chronological order has been indicated by the letters A, B, and C, whereas the ranks are specified 1st, 2nd, or 3rd. Hence, Observation C occurred after Observation B, but would be ranked before it, because relative values along the response dimension define ranks, not the temporal order of occurrence.

The equation above can be rearranged to specify the rank that must serve as the criterion ($k$) in order that criterional events occur with a specified probability $w$. If $w = 1 - [k/(m + 1)]$, then subtracting 1 from both sides and reversing signs yields $1 - w = k/(m + 1)$, and multiplying both sides of the equation by $m + 1$ yields $k = (m + 1)(1 - w)$. This expression, the basic percentile equation, provides the criterion rank order $k$ that must be exceeded to witness a criterional response with a trainer-specified probability $w$, given a sample of $m$ prior observations. For example, to observe a criterion response half the time (i.e., $w = .5$) given three prior observations (i.e., $m = 3$), the rank order to be exceeded would be set to $k = (3 + 1)(1 - .5) = 2$. A response would be considered criterional if it exceeded the value represented by the 2nd rank (C) in the middle panel of Figure 4.

By specifying a criterion relative to a distribution of behavior rather than relative to some fixed physical value, the probability of a criterional response on the next trial is determined not by the client, but by the trainer. But what of subsequent trials? That is, suppose the next response occurs at the value denoted by D in the bottom panel of Figure 4. How should the criterion for the next response be evaluated? One possibility would be to add each new observation to the list of previous values, increasing the value of $m$ in the above equation. However, this strategy does not distinguish current observations from more remote ones. Although above I noted that chronological order was unimportant in ranking observations, we must at some point recognize that as observations become more and more remote, they may no longer adequately characterize the population of values likely to be observed now. An alternative that overcomes this problem is to update the distribution of observations by replacing the oldest observation in the distribution (A) with the newest (D). Doing so maintains the number of comparison observations ($m$) constant, such that with a constant $w$ the required value of $k$ is unaltered. For example, to observe a criterion response with a probability of .5, it is again necessary that the next response exceed two of the three observations. This would now mean that the criterion would be set at the value represented by D, because the comparison distribution now consists of Observations B, C, and D, with the latter ranked 2nd of the three (see the bottom of Figure 4). Note that the current response becomes part of the comparison memory only after all decisions concerning its criterional status have been made. The current response is never compared to itself, given that, by definition, it can never exceed itself.

Updating the comparison distribution with each response explicitly does what skilled shapers do. Although we start differentiating responding where the subject is, as the comparison distribution changes, the criterion changes with it to keep the overall probability of a criterion response relatively constant. Percentile schedules set the criterion at a fixed rank to control overall criterional response probability, and then let changes in the physical values of the responses comprising the comparison distribution determine where the physical criterion will fall. Note that only the physical value of the criterion changes; the rank order (i.e., $k$) of that physical value remains constant. Hence, behavior is differentiated only in the sense that physical values change; the criterion remains constant at a particular rank order, providing a constant probability of criterional responses. In this way, percentile schedules concurrently shape behavior (changing values along a physical dimension) while controlling the probability of criterional responses (defined ordinally).

At this point, it may be helpful to consider a hypothetical example of how a percentile schedule might be programmed in an applied setting. In posing any particular example, I ask the reader's indulgence in granting me a degree of latitude sufficient to illustrate the workings of the schedule without tackling the myriad additional problems that make applied work especially challenging, but that are not uniquely relevant to the use of percentile schedules (e.g., defining responses, observing and recording strategies, patient compliance, etc.). Failure to mention these important aspects of behavior analysis should not be taken as an admission of ignorance or disregard. They represent complex problems that must always be addressed, independent of the exact procedure used, and for this reason are not germane to a discussion of percentile schedules per se.

To illustrate how a percentile schedule might be used to help define response criteria in an applied setting, consider an example involving "task engagement." Assume I wish to increase the time a developmentally disabled client devotes to a vocational task while at the work station. I might

| Minute | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Baseline | Pr=.2 | | | | | | | | | | | | | | | | | | | |
| Percentile | | | | | | | | *w*=.2 and *m*=4 | | | | | | | | | | | | |
| Current criterion | - | - | - | - | - | - | - | 4 | 4 | 5 | 5 | 5 | 5 | 4 | 3 | 3 | 5 | 5 | 8 | 8 |
| On-task intervals | 1 | 3 | 2 | 3 | 4 | 2 | 1 | 2 | 5 | 4 | 3 | 2 | 3 | 3 | 3 | 5 | 2 | 8 | 6 | 4 |
| Reinforcement | | | | | | | | | | | | | | | | | | | | |
| Comparison Distribution | | | | ↑ | ↑ | ↑ | ↑ | Current Criterion | | | | | | | | | | | | |
| During Minute: 8 | | | | 3 | 4 | 2 | 1 | 4 | | | | | | | | | | | | |
| 9 | | | | | 4 | 2 | 1 | 2 | 4 | | | | | | | | | | | |
| 10 | | | | | | 2 | 1 | 2 | 5 | 5 | | | | | | | | | | |
| 11 | | | | | | | 1 | 2 | 5 | 4 | 5 | | | | | | | | | |
| 12 | | | | | | | | 2 | 5 | 4 | 3 | 5 | | | | | | | | |
| 13 | | | | | | | | | 5 | 4 | 3 | 2 | 5 | | | | | | | |
| 14 | | | | | | | | | | 4 | 3 | 2 | 3 | 4 | | | | | | |
| 15 | | | | | | | | | | | 3 | 2 | 3 | 3 | 3 | | | | | |
| 16 | | | | | | | | | | | | 2 | 3 | 3 | 3 | 3 | | | | |
| 17 | | | | | | | | | | | | | 3 | 3 | 3 | 5 | 5 | | | |
| 18 | | | | | | | | | | | | | | 3 | 3 | 5 | 2 | 5 | | |
| 19 | | | | | | | | | | | | | | | 3 | 5 | 2 | 8 | 8 | |
| 20 | | | | | | | | | | | | | | | | 5 | 2 | 8 | 6 | 8 |

Figure 5. Hypothetical example illustrating the operation of percentile schedules.

decide to divide each minute into 5-s intervals and record whether the client remained "on task" during the entire interval. (Other recording techniques might be preferable; this example has been chosen to provide a moderate range of observable values.) This regimen generates a score each minute ranging between 0 (no intervals on task) and 12 (consistently on task). Presume further that a reinforcer has previously been identified for this client. The first decision required is how often to provide reinforcement. Assume that my previous history with this client suggests that, on average, a reinforcer delivered every 5 min should suffice.

Although technically not required, a baseline condition would probably be conducted first. This serves two functions: (a) It provides an indication of the current level of behavior, and (b) it provides a comparison for the effects of intermittent reinforcement alone. The baseline provides the same reinforcement frequency as the percentile procedure but does not make it contingent on the degree of task engagement. This can be accomplished by presenting a reinforcer at the end of each 1-min interval with a random probability of .2 (i.e., on the average, every 5th min).

Suppose this procedure generates the behavior shown for the first 7 min in Figure 5. The row labeled "current criterion" contains no value under baseline, because it is by definition undefined (i.e., reinforcement is not contingent on task engagement). Reinforcement (denoted by the darkened square in the row labeled "reinforcement") is delivered by the prearranged random probability after Minutes 2 and 6. The mean number of on-task intervals during baseline is approximately 2.3 per minute, and the mean number of intervals engaged prior to reinforcer delivery is 2.5. The correspondence between these values indicates that reinforcement is nondifferential (i.e., the degree of task engagement does not influence reinforcement delivery).

Instituting a percentile procedure involves first substituting .2 for $w$ in the equation above, such that $k = (m + 1)(1 - .2) = .8(m + 1)$. Next, a value must be assigned to $m$ (i.e., how many prior observations will be used as comparison values?). This is not a completely arbitrary decision, as will be discussed shortly. For purposes of the current example, however, setting $m$ to 4 yields a value for $k = .8(5) = 4$. Hence, to observe criterional responses with a probability equal to .2 ($w$), the current score should be compared to the score from the most recent four intervals ($m$), and if it exceeds the fourth rank-ordered value ($k$) in the list of prior scores, it is criterional. If criterional and only criterional responses are reinforced, the probability of reinforcement and the probability of criterional responses will be equal (i.e., .2).

The most recent four scores during Minute 8 are those from Minutes 4 through 7. As such, they define the initial comparison distribution for the percentile condition. During that time, my client was on task for three, four, two, and one interval(s), respectively. Ordering these observations from lowest to highest associates one, two, three, and four intervals with Ranks 1 through 4, respectively. Hence, the current criterion score is 4 (the 4th-ranked value). The score during Minute 8 is actually 2. Because this does not meet the criterion, no reinforcement is delivered. For the next minute, the comparison distribution is changed by replacing the most remote score with the most recent one. The comparison distribution now comes from Minutes 5 through 8, or four, two, one, and two intervals on task, which when ordered becomes 1, 2, 2, and 4. Thus, the criterion score remains 4. During this minute, the number of intervals engaged in the task (five) exceeds the criterion and reinforcement is delivered. During the next 4 min, the 4th-ranked value in the comparison distribution is always five, the current score never exceeds the criterion, and no reinforcement is provided. In Minute 14, the 4th rank in the distribution falls back to 4, and in the next minute it falls back to 3, because the current distribution now contains three, two, three, and three intervals engaged. This illustrates how percentile schedules, while maintaining contingencies to differentiate "more" engagement (i.e., the largest value in the comparison distribution sets the criterion), also allow the definition of "more" to slide back towards a lower score if behavior consistently moves in that direction. During Minute 15, the number of engaged intervals equals the criterion, raising the question of how to treat ties. Classifying all ties as criterional overestimates the expected probability, whereas classifying them as noncriterional underestimates it. The problem is magnified as ties become more frequent. If a sequence of 20 observations all tied, treating all as criterional or noncriterional would result in a reinforcement probability of 1 or 0, respectively, across those observations, and not the .2 probability programmed. The simplest solution is to select ties with a random probability equal to $w$ and call them criterional. Here, the observation is not classified as criterional and as a result does not generate a reinforcer (that this is a function of a tie decision is indicated by the gray shading in the figure). Continuing through Minute 20 results in two more criterional scores, a score of 5 in Minute 16, when the criterion was 3, and eight intervals engaged 2 min later when the criterion was 5. As a result, three criterional responses are observed during the 13 min of the percentile procedure, for an obtained probability of .23, within the sampling error of the probability programmed. The mean score during this time was 3.8, whereas the mean score preceding reinforcement was 6.0. This difference defines the differential nature of the reinforcement contingency.

The above equation defining the basic percentile schedule provides a fixed, specified probability of a criterional response at all times during the course of shaping. If reinforcers follow only criterional responses, as is most likely in applied settings, the probability of reinforcement will also equal $w$, and reinforcement will be maximally differential. That is, reinforcers will be delivered only after responses that are relatively closer to the terminal response along the dimension being differentiated (e.g., "more" time on task). Criteria are always set relative to current behavior, so not only does training start with the client, it stays and ends with the client. Criteria are updated with each response, remaining most responsive to changes in behavior. Finally, the probability of a criterional response can be specified to be whatever works best for that particular client–response–reinforcer combination, maximizing the differential nature of the contingency (i.e., reinforcing only instances of "more" engagement, as defined by the current distribution) but providing reinforcement consistently and intermittently in order to maximize persistence and decrease the probability of the frustration, aggression, emotional responses, and response elimination associated with extinction. Determining the optimal criterional response probability is not a problem unique to percentile schedules, but rather is an empirical question faced with each new procedure and/or response. Unlike other procedures, percentile schedules allow an empirical answer to this question by directly controlling reinforcement parameters independent of behavior.

Percentile schedules appear to meet all the requirements for a viable procedure to formalize shaping except the last—they do not specify a terminal response. The criterion is never specified as an absolute; rather, it is described only in relative fashion (i.e., exceed the $k$th rank). Yet, even here, percentile schedules help to focus our understanding of shaping. There is only one terminal response of all shaping—to do better on the next trial than on previous trials. This is what percentile schedules program, where "better" is defined as exceeding the $k$th rank and "previous trials" is given by the most recent $m$ observations. Because criteria are evaluated relative to ongoing behavior, there is never a need to stop shaping. Once acquisition is complete and behavior stabilizes, the percentile schedule can still be used to select the same overall proportion of responses for reinforcement. And if an external event should disrupt responding, the criteria automatically adjust to take this into account and shape responding back to previous levels while ensuring a constant reinforcement density.

### *Implementing Percentile Schedules in Applied Settings*

Percentile schedules make two requirements not found in traditional operant conditioning procedures. First, they require a continuously updated record of the most recent $m$ responses that, second, can at least be partially ranked. Modern computers of all sizes are fast enough to program percentile contingencies with no discernible delays between responses and reinforcers. But computers are not a prerequisite to the implementation of percentile contingencies. Selecting particular pairs of criterional response probabilities and comparison distribution sizes makes it possible to program percentiles with a pencil and paper.

An updated list of the last $m$ responses can easily be kept on a device similar to that shown in Figure 6. A roll of paper like that used in event recorders, adding machines, and cardiac monitors is threaded through a window with lines demarcating each comparison observation. Each new response is recorded in a slot marked "current response," located just below the window. Once recorded, the paper is advanced one observation out the top of the device, such that the latest observation (at the bottom) replaces the oldest one (at the top) in the list. An adjustable shutter slides vertically to vary the number of prior observations ($m$) visible, as noted by the numbers along the right side. In this way, the most recent $m$ responses will always show in the window to provide the current comparison memory.

Ranking the observations is a more difficult, but surmountable, task. Even in the laboratory we do not spend the time required to rank all observations and then determine where the physical value of rank $k$ is. Rather, the computer code compares the current value to every value in the comparison distribution. If the current value exceeds the comparison value, a counter is incremented. As soon as that counter exceeds $k$, the response is considered to be criterional. If, after all comparisons have been made, the counter still does not exceed $k$, the response is noncriterional. This procedure is faster than sorting, but is still of little help to someone lacking a computer.

An alternative approach is illustrated in the top panel of Figure 4. The expected probability that the next response will fall into each interval is given by $1/(m + 1)$. Although it is difficult to sort through a whole list of values, it is relatively easy to scan a list and determine the largest value. If we always set the criterion to the largest value in the comparison distribution, criterional responses will be observed with a probability equal to $1/(m + 1)$. To observe a particular probability $w$ of a criterional response, therefore, $m$ can be set to the value given by $m = (1/w) - 1$. For example, to obtain a criterion probability of .2, $m = (1/.2) - 1 = 5 - 1 = 4$ observations are needed. This was the approach used in the task-engagement example above. The suggestion of the sliding window in the above device now may be more understandable. By using this calculation to determine the memory size, and then adjusting the window to the appropriate number of values, it is necessary only to determine whether the current value is longer than any currently in view. If so, it is a criterional response, and will occur with the probability given by $w$.

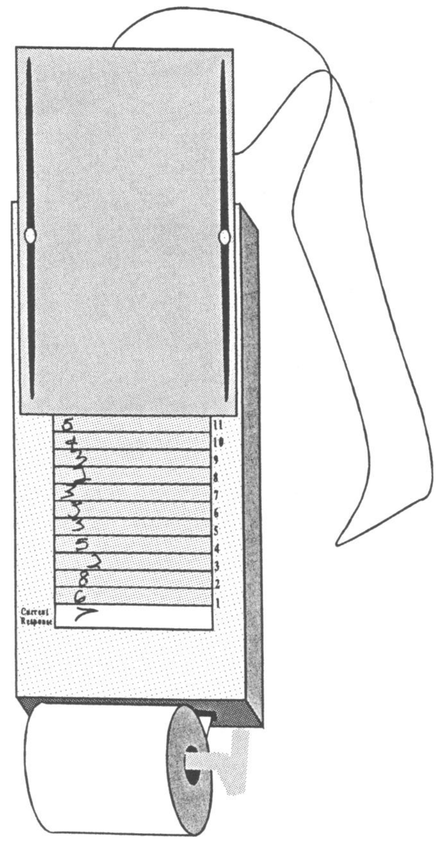


Figure 6. Illustration of a recording device designed to facilitate the programming of percentile contingencies in the absence of computers. Response values are recorded on the strip of paper in the slot at the bottom of the window. Prior to each new trial or response, the paper is advanced one response up, such that only the most recent responses are visible. The vertical shutter can be used to adjust the number of previous observations visible at any time ($m$). See text for further description of use.

### *Limitations on Applying Percentile Procedures*

There are only two formal assumptions involved in deriving the percentile schedule equation: (a) Behavior must be measured in such a way that ordinal ranks can be assigned, and (b) those ranks must not be sequentially related (i.e., successive observations must represent random and independent samples from the population of response values). Dealing with the second limitation first, the degree to which responses are sequentially related increasingly affects the operation of percentile schedules as the comparison distribution size ($m$) decreases. Consider a situation that violates the assumption that observations occur randomly and independently. Suppose that each response has a .8 probability of being followed by an even longer value (i.e., four out of every five times, the next value is longer). Suppose further that the programmed probability of a criterional response ($w$) is .5. Using the shortcut suggested above, this probability is plugged into the equation for $m$ to solve for the memory size that will yield the appropriate

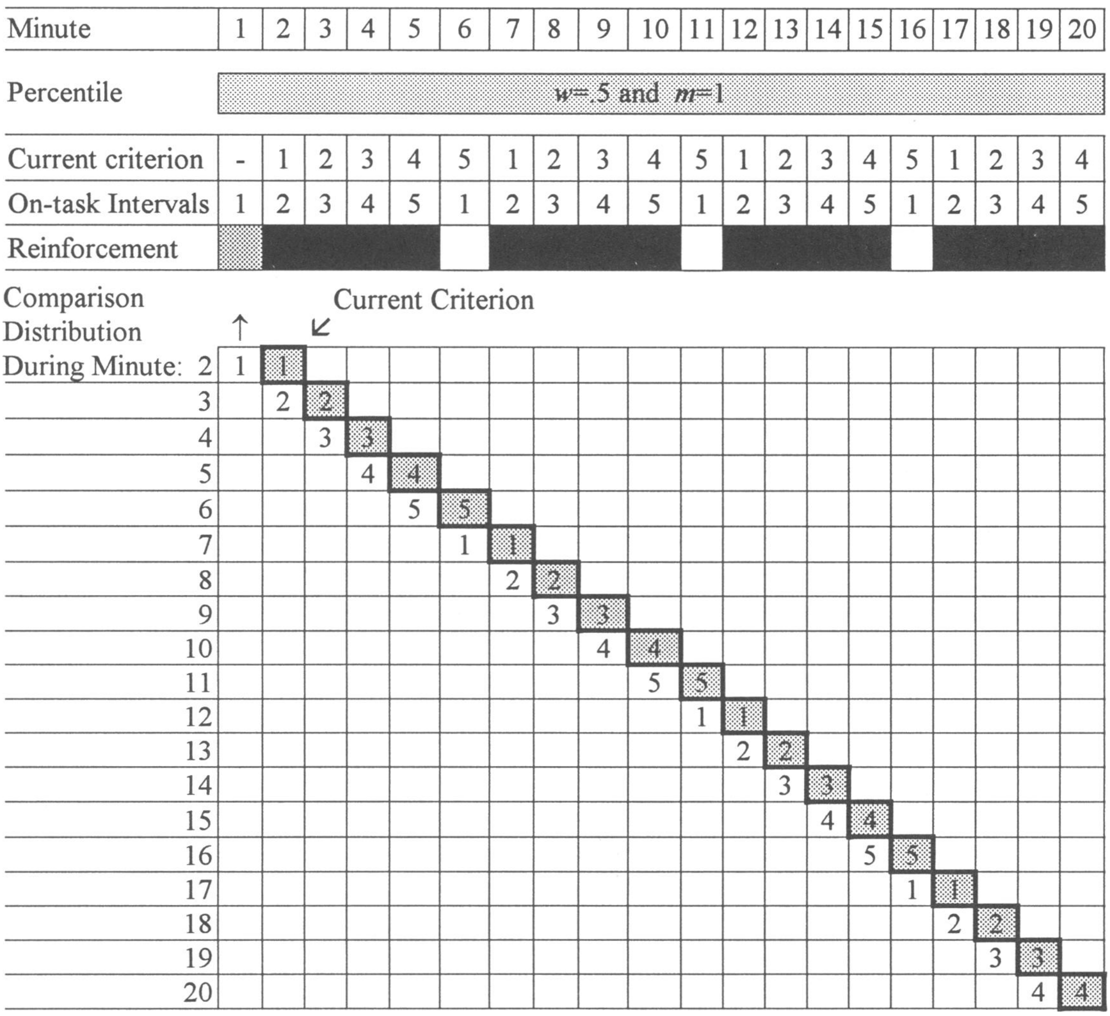


Figure 7. Effects of nonrandom observations on a percentile schedule with $w = .5$ and $m = 1$.

probability if the current value exceeds the longest comparison observation. This value is $m = (1/w) - 1 = 1/.5 - 1 = 1$. Hence, using only the most recent response as a comparison value, we begin programming the percentile schedule.

Figure 7 presents an illustrated case in which $m = 1$ and $w = .5$. To satisfy the requirement of a .8 probability of a longer subsequent observation, the score (using the task-engagement example again) repeats a cycle of values 1, 2, 3, 4, and 5. Because the first value observed has no comparison, suppose we arbitrarily call it criterional with a probability of .5. This is denoted in the reinforcement square by gray shading. During Minute 2, the criterion is one and two intervals are scored as engaged, so the response meets criterion and is reinforced, as indicated by the dark square in the row labeled "reinforcement." In the next minute the criterion is two and the score is 3, so again reinforcement is delivered. The criterion in the next 2 min is three and four, respectively, and the number of intervals engaged is four and five, so reinforcement is provided on each occasion. In Minute 6, when the cycle begins again, the criterion is five but only one on-task interval is scored, so reinforcement is not delivered. During Minutes 2 through 6, four of five scores are considered criterional and are reinforced. This represents a substantial departure from the .5 probability nominally programmed by the percentile schedule. Further, there is no indication

| Minute | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Percentile | w=.5 and m=3 | | | | | | | | | | | | | | | | | | | |
| Current criterion | - | - | - | 2 | 3 | 4 | 4 | 2 | 2 | 3 | 4 | 4 | 2 | 2 | 3 | 4 | 4 | 2 | 2 | 3 |
| On-task Intervals | 1 | 2 | 3 | 4 | 5 | 1 | 2 | 3 | 4 | 5 | 1 | 2 | 3 | 4 | 5 | 1 | 2 | 3 | 4 | 5 |
| Reinforcement | | | | | | | | | | | | | | | | | | | | |

Comparison Distribution — Current Criterion

| During Minute: | | | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 4 | 1 | 2 | 3 | 2 | | | | | | | | | | | | | | | | |
| 5 | | 2 | 3 | 4 | 3 | | | | | | | | | | | | | | | |
| 6 | | | 3 | 4 | 5 | 4 | | | | | | | | | | | | | | |
| 7 | | | | 4 | 5 | 1 | 4 | | | | | | | | | | | | | |
| 8 | | | | | 5 | 1 | 2 | 2 | | | | | | | | | | | | |
| 9 | | | | | | 1 | 2 | 3 | 2 | | | | | | | | | | | |
| 10 | | | | | | | 2 | 3 | 4 | 3 | | | | | | | | | | |
| 11 | | | | | | | | 3 | 4 | 5 | 4 | | | | | | | | | |
| 12 | | | | | | | | | 4 | 5 | 1 | 4 | | | | | | | | |
| 13 | | | | | | | | | | 5 | 1 | 2 | 2 | | | | | | | |
| 14 | | | | | | | | | | | 1 | 2 | 3 | 2 | | | | | | |
| 15 | | | | | | | | | | | | 2 | 3 | 4 | 3 | | | | | |
| 16 | | | | | | | | | | | | | 3 | 4 | 5 | 4 | | | | |
| 17 | | | | | | | | | | | | | | 4 | 5 | 1 | 4 | | | |
| 18 | | | | | | | | | | | | | | | 5 | 1 | 2 | 2 | | |
| 19 | | | | | | | | | | | | | | | | 1 | 2 | 3 | 2 | |
| 20 | | | | | | | | | | | | | | | | | 2 | 3 | 4 | 3 |

Figure 8. Effects of nonrandom observations on a percentile schedule with $w = .5$ and $m = 3$.

that this probability will decrease with further sampling—each time the cycle repeats, four out of every five scores are considered criterional. This is because percentile schedules are based on the presumption that the next observation is as likely to be ranked lowest as highest—that each interval has an equal likelihood of claiming the next response. Yet in this example, responses are falling on the high side of the last observation 80% of the time. This asymmetry defines a sequential dependency. Knowing the most recent response allows prediction of the next one's value (i.e., 80% of the time the next response will be longer than the current one, rather than the expected 50% that it will fall either above or below).

Although sequential dependencies diminish the ability of percentile schedules to control criterional response probability, their effects can be minimized by increasing the comparison distribution size. Consider the same scenario, but with a memory size of three. From the percentile equation above, $k = (m + 1)(1 - w) = (4)(.5) = 2$; the current response will be considered criterional if it exceeds two of the three most recent values. Figure 8 illustrates this case. Reinforcement is randomly assigned with a probability of .5 after each of the first 3 min, because of the lack of sufficient comparison observations. In Minute 4, the criterion is two and the score obtained is 4, so reinforcement is delivered. In the next minute the criterion is three and the client remains on task for five intervals. During both Minutes 6 and 7 the criterion is four, but only one or two intervals are judged as being on task; hence, no reinforcement occurs. In the next minute the criterion is two, and three on-task intervals are scored, resulting in reinforcement. The cycle then repeats. Adding two additional observations to the comparison distribution increases the correspondence between programmed and obtained criterional response probability over that generated with only a single comparison value. Where the obtained probability of criterional responses had been .8 when $m = 1$, here it is only .6, which is much closer to the nominally programmed value of .5. Task engagement continues to show strong sequential dependencies, in that scores during 4 out of every 5 min are longer than their immediate predecessor. However, these scores are no longer guaranteed to be criterional because it is not sufficient to exceed only the last score; rather, the current score must exceed two of the three most recent responses. Because more than a single score is used as a comparison, the effects of local variation, and hence sequential dependencies, are diminished.

Although not always eliminated, the effects of sequential dependencies are greatly limited by increasing the comparison distribution size. This property allows percentile schedules to be used even in situations in which the formal assumptions underlying their derivation are violated. However, their use in these circumstances would likely be limited to those involving computer control, because comparisons would have to be made to more than the largest comparison value.

The other "limitation," that responding be ordinally rankable, could actually aid application of percentile schedules. In the laboratory, percentiles have been used exclusively to shape responding along a single dimension—to shape longer or shorter interresponse times (e.g., Alleman & Platt, 1973; Arbuckle & Lattal, 1992; Galbicka & Platt, 1986; Kuch & Platt, 1976), response durations (e.g., Platt, 1984; Platt, Kuch, & Bitgood, 1973), run lengths (e.g., Galbicka, Fowler, & Ritch, 1991; Galbicka, Kautz, & Jagers, 1993; Galbicka, Kautz, & Ritch, 1992), different spatial response locations (e.g., Davis & Platt, 1983; Galbicka & Platt, 1989; Scott & Platt, 1985), or variable response sequences (Machado, 1989). Although there may be applications that could easily be envisioned as changing a single dimension of behavior (e.g., the running or task-engagement regimens described earlier), it is more likely that behaviors to be shaped in both research and applications with human subjects will involve behavioral change along multiple dimensions. Although to this point changes along single dimensions have been emphasized, the fact that percentile schedules do not deal directly with physical aspects of behavior, but only their ranking relative to prior behavior, implies that the number of response dimensions involved is immaterial. Any behavioral sequence, from any identified starting point to any terminal response, can be subjected to the workings of a percentile schedule to the extent that each response in the sequence can reasonably be ranked relative to every other response. This is not a burden unique to percentile schedules because all shaping requires some means of attributing directional change (and hence some crude ranking) to behavior across many dimensions. Percentile schedules actually remove some of the burden by identifying steps that require a modified response definition.

To illustrate, suppose we wish to train a developmentally disabled client to drink fluid though a straw. Prior observation of the behavior leads the shaper to suggest that the following five behaviors might be involved: (1) holds glass, (2) directs glass toward mouth, (3) holds straw with other hand, (4) directs straw into mouth, and (5) sucks on straw. These five behaviors can easily be ranked 1 to 5, with 1 being furthest from the terminal response and 5 being closest. A percentile schedule could be imposed by recording the response value (i.e., 1 through 5) on each trial. Whether our conception of the response matches the subject's will be evident in the relative frequency of each of the different rankings. For example, if after collecting substantial data an analysis revealed that Step 3 (i.e., holds straw with other hand) was never recorded, but instead the client progressed directly from Step 2 to Step 4 by holding the glass and manipulating it to position the straw in his or her mouth, then Step 3 probably represents an overly specific or functionally irrelevant response category for this subject. On the other hand, if some other step showed a disproportionately high frequency of occurrence, it is possible that further definition may be required to distinguish functionally the multiple response classes that are likely being included under this heading. The interesting thing about percentile schedules is that making such modifications on line does not alter their effectiveness. Refining response definitions by adding or removing particular classes at any time during shaping has no effect because the *number* of response classes is irrelevant to the operation of percentile schedules; the only requirement is that these response classes must easily be ordinally rankable with respect to all other responses. By allowing behavior to be the ultimate arbitrator in deciding where response classes "naturally fracture," the potential ramifications of misjudging these classes are limited. More classes can be added as they are identified with experience, and superfluous ones can either be dropped or retained.

*The Shape of Things to Come?*

I have argued that percentile schedules represent a formalization of the rules of shaping. I have tried to present the derivation of the equations that define percentile schedules in such a way as to make this relation clear. Finally, I have tried to demonstrate that little instrumentation is necessary to program these procedures. The applied utility of percentile schedules ultimately rests in the hands of the applied community. I only wish to note that percentile schedules do not impose any additional constraints and may actually remove some of those that limit application of other behavioral procedures. The ever-present problems of observation and measurement in applied settings will obviously affect application of percentile schedules as well. However, because percentile schedules use only ordinal response values, many problems associated with traditional procedures may be circumvented.

The research potential of these procedures is equally far-reaching. I have treated the probability of criterional responses and of reinforcement as equivalent throughout most of this paper, but this is true only if all criterional and only criterional responses are reinforced. This is generally the case in applied settings because these conditions most rapidly change behavior. However, a research setting holds one further level of complication. The more general scenario characterizing all attempts at response differentiation is represented in Figure 9. Training involves two independent decisions on the part of the experimenter: (a) whether the response is criterional or not, as indicated by whether it falls into rows indicated W and $\overline{W}$, and (b) subsequently whether or not to reinforce that response. This second decision is independent of the first and determines whether the response falls into the two columns labeled Z and $\overline{Z}$. Responses falling in Cell *a* are criterional responses that were followed by reinforcement, and those in Cell *b* are criterional responses that were not reinforced. Similarly, responses in Cell *c* are noncriterional reinforced responses, and those in Cell *d* are noncriterional nonreinforced responses. These cells describe all potential joint occurrences of responding and reinforcement. The percentile equation specifies segregation by criterional response probability only (i.e., it specifies whether the current response falls in the upper row or the lower one); it technically does not speak to whether that response, criterional or not, is reinforced.

Two additional parameters are needed to define the conditional probability of reinforcement for criterional and noncriterional responses, termed $u$ and $v$, respectively (e.g., Galbicka & Platt, 1986, 1989; Scott & Platt, 1985). These specify the probability, given the prior occurrence of a criterional (or noncriterional) response, that a consequence also occurs. As the conditional probabilities of reinforcement increase, the relative proportion of responses on the left side of the appropriate row increases. Thus, increasing $u$ increases the proportion of criterional responses falling into Cell *a* (i.e., without changing the total number of responses in the top row, increasing $u$ shifts responses from Cell *b* to Cell *a*). Increasing $v$ increases the proportion of noncriterional (bottom row) responses falling into Cell *c*. Each of these probabilities is independently specifiable, from each other and from the probability of a criterion response ($w$).

The matrix in Figure 9 represents an extensive

**Consequence**

| Criterional Response | | $Z$ | $\bar{Z}$ | |
|---|---|---|---|---|
| | $W$ | $a$ | $b$ | $u=\Pr(S^R \mid W)=a/(a+b)$ |
| | $\bar{W}$ | $c$ | $d$ | $v=\Pr(S^R \mid \bar{W})=c/(c+d)$ |

$$w=\Pr(W)=(a+b)/(a+b+c+d)$$

Figure 9. Two-by-two table showing the joint effects of specifying the probability of a criterional response ($w$) and the conditional probabilities of reinforcement for criterional ($u$) and noncriterional ($v$) responses. See text for details.

research program. By using the percentile equation to generate a specific probability $w$ of a criterional response, the experimenter gains control over how responses get segregated into the two rows of the matrix. Specifying particular values of $u$ and $v$ allows researchers to specify further how events in each of those two rows are distributed with respect to reinforcement. Combining these two operations for the first time allows researchers to dictate *in advance* the expected frequencies of each of the four cells, and hence to control the relation between criterional responding and reinforcement. This prompts a number of previously unanswerable questions. For example, what happens to shaping as $u$ increases? as $v$ increases? as $w$ increases? In general, we might expect that increasing $u$ (reinforcement of criterional responses) would facilitate shaping, whereas increasing $v$ (reinforcement of noncriterional responses) would impede it. But what of the effect of $w$? The answer here is less clear. The suggestion offered by Rule 3 is that optimal shaping will occur when the probability of a criterional response is not so small that contingent reinforcement is too infrequent, but not so large as to be virtually nondifferential. Hence, modifying $w$ may produce an inverted U-shaped function with respect to response differentiation. How might this interact with the relations above? That is, given that increasing $u$ enhances acquisition and increasing $v$ retards it, how are these effects quantitatively modified by the proportion of responses considered criterional? In other words, what happens to the slope of these functions as $w$ changes? Finally, is there a way we can put all these effects together into some general index that will correlate with the degree of shaping under a particular set of parameters? For example, if increasing $u$ and decreasing $v$ each enhance acquisition, then acquisition is differentially correlated with increasing values in Cells $a$ and $d$ relative to Cells $b$ and $c$ in Figure 9. Several measures of statistical contingency also increase with such changes and might be used as metrics of the relation between reinforcement and criterional responses that map to the changes in behavior produced (see Galbicka & Platt, 1986, 1989; Scott & Platt, 1985, for examples of such efforts involving nonhuman subjects).

Percentile schedules represent a radical departure from traditional methods in the degree of experimental control they afford. They represent the next step in the experimental control of variables that are critical to the modification of behavior; thus, they could potentially revolutionize our methods as well as our models of behavior change. Specific examples of applications presented here can serve only to focus attention unduly on particular behaviors or populations in the applied realm, at the expense of other, potentially more relevant behaviors or populations. But with that caveat in mind, percentile schedules might, for example, provide for replacement of vague qualitative labels for "cognitive abilities" with more extensive, quantitative formulations of learning and learning disabilities. A diagnostic video game could be devised that repeatedly requires the client to select from a ring of keys the one that unlocks the door barring entrance into the next passage. The keys differ along a number of dimensions, but the one of interest is key length. Selecting a particular key is considered criterional if it is closer to some target length, chosen by the experimenter, than 80% of the previous $m$ key selections. This is an example of a targeted percentile in which responding is shaped to a particular value instead of in a particular direction. Targeted percentile schedules are programmed as two separate percentile schedules. One operates when the current response is short of the target (i.e., when the key selected is less than the target value) and differentially reinforces values *longer* than the 80th percentile of all recent previous selections that were also short of the target. The other operates when the current response is above the target and differentially reinforces responses *shorter* than 80% of all responses above the target. These two, simultaneously operating percentile schedules control the overall probability of criterional responses above and below the target value, but reinforce only the closest 20% of all responses, causing the distribution to peak in the area of the target (see Galbicka et al., 1993, for an example of targeted percentile shaping).

The rate of acquisition of this task under a variety of different values of $w$, $u$, and $v$ could be used to increase the precision of behavioral "diagnoses" of different populations. For example, certain de-

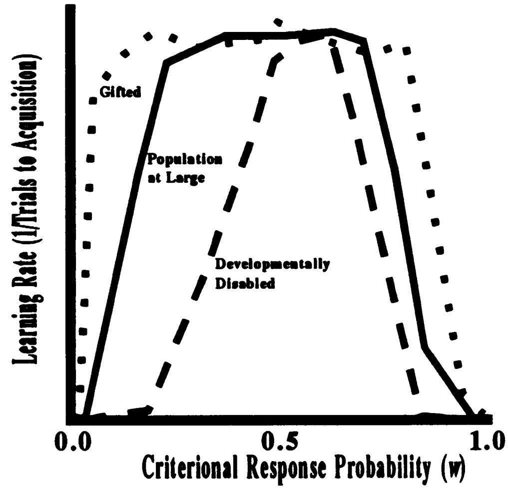


Figure 10. Hypothetical results showing the rate of learning (1/number of trials to acquisition) as a function of manipulating the probability of a criterional response ($w$) under a percentile schedule. The solid curve shows potential results from the population at large, the dashed curve illustrates potential results from a population of developmentally disabled clients, and the dotted line depicts results from a population of gifted individuals.

velopmental disabilities might be associated with generally slower acquisition than other populations, under all values of $w$, when $u$ and $v$ are constant at 1.0 and 0. Or the relation may be more complex, such as that shown in Figure 10. The possibility exists that the degree of learning demonstrated under different $w$ values by different populations could be used to categorize a range of behavioral populations. The solid curve suggests that, in general, the population at large may shape well under a relatively wide range of $w$ values, with learning falling off at the extremes when $w$ is relatively low, and hence reinforcement density is low, or when it is relatively high; consequently, criterional responses can be diverse at any particular time. Developmentally disabled clients might demonstrate a more restricted range of effective parameters (shown by the dashed line) indicating a greater sensitivity to sparse reinforcement (i.e., low $w$ values) and a greater difficulty coming under control of the dimension of responding being shaped at higher $w$ values. On the other hand, gifted populations might be defined by the production of functions similar to those shown by the dotted line, with learning spread across a range of $w$ values broader than that for the population at large.

These hypothetical examples are presented only to illustrate the potential quantitative analysis of learning that could be derived using percentile procedures. Percentile schedules provide a better platform from which to develop such "diagnostics," because they greatly reduce differences arising from the different behavioral repertoires each subject brings to the task. Traditional operant procedures specify behaviors with particular physical characteristics (responses with particularly defined topographies, forces, durations, etc.). The variable nature of operant behavior, in conjunction with the variable histories of reinforcement each subject brings to the training environment, all but guarantee that different subjects will respond with different frequencies of responses meeting the fixed physical criteria. As a result, each subject's behavior will interact uniquely with the contingencies of reinforcement. Percentile schedules remove the obstacle posed by individual behavioral differences by defining all response criteria relative to each subject's own repertoire. All subjects share exactly the same probability of emitting a response closer than half the responses emitted recently, independent of how close those recent responses were. This holds not only across subjects but also across time within the same subject.

In addition to providing a potential diagnostic tool, percentiles could be used to assess the effects of variables such as alcohol consumption and drug administration, sleep deprivation, or aging. They may also serve to categorize the degree of behavioral disruption induced by reinforcement for noncriterional responses (i.e., "distraction") or by lack of reinforcement for criterional responses (i.e., "short attention span"). Not only could these be used as diagnostic tools, they also could serve as therapeutic aids by providing a consistent probability of reinforcement while encouraging continued development towards some targeted value.

Finally, percentile schedules may help revive interest in programmed texts (e.g., Holland & Skinner, 1961) and associated personalized systems of instruction. Programmed texts enjoyed a certain popularity for a period of time in the 1960s and 1970s, but have recently declined in popularity. Students often complained that these texts were boring. Because step size could not be individualized in a printed text, it was often set at a size almost everyone could achieve, but one most students found extremely small.

The advent of computer-based instruction (e.g., McDade & Goggans, 1993) provides a means of presenting varying stimulus frame sequences, and percentile schedules could provide a mechanism for coupling that presentation to performance. Numerous multiple choice questions could be developed to illustrate each key concept in successive organizational units of the program. Units of the program would be ranked relative to other units to establish a progression through the course. The lowest ranked unit would present questions on fundamental concepts, shaping a transition from introductory or lay concepts to those required by the subject matter. Subsequent units would be ranked relative to the degree to which correct responses in this unit depend on the acquisition of textual behavior in some prior unit. Each correct response would generate a rank appropriate to the level of that question, and achieving some number of consecutive correct questions at one level would be prerequisite to moving on to the next. Incorrect answers would generate ranks associated with the prior unit in which material relevant to that answer was first presented. For example, in a text on behavior analysis, a question dealing with an example of negative reinforcement might include an answer that makes reference to a decrease in response rate. Selecting this option would constitute a lower ranked response because the more fundamental verbal operant that reinforcement is always associated with an increase in the frequency of an operant is not evident in the repertoire. By ranking this response with the level in which the general concept of reinforcement was introduced, this error would not only force return to a prior level, it would specifically return to the level associated with the misunderstanding leading to the inappropriate selection of that answer. Specific deficits could then be corrected while adjusting the criterion automatically to generate a relatively constant frequency of reinforcement.

The application of percentile procedures may prove to be difficult in the extreme, or the definition and transduction of multidimensional response classes too cumbersome to be practically useful. Even so, the present discussion still serves the heuristic function of explicitly dissecting the fundamental components of shaping. Increasing familiarity with percentile schedules leads to the recognition that a variety of contingencies in daily life are arranged, not necessarily with any conscious awareness, in much the same way as percentile schedules. As social organisms selected to select behavior, we often act as organic ranking machines. We maintain continuously updated lists of our companions' recent accomplishments and provide reinforcement with respect to some relatively constant upper proportion of this list. We judge our children not with respect to some fixed standard, but relative to their development (i.e., the degree of change in the behavioral distribution). Once responses become commonplace, they no longer merit reinforcement. Rather, the relatively rarer exceptional response that may yet be quite far from target but still qualifies as criterional because it is *closer* wins our attention. As people mature, behavior changes but the criterion remains qualitatively the same—do better, be in the upper tail of the distribution, be exceptional, relative to what you have been recently. When those around us get sick or old, we relax the criterion by increasing the probability of a criterional response or more often by noting that the current distribution of responding has regressed.

Understanding percentile schedules may increase our understanding of the complex social and nonsocial dynamics that shape behavior. They allow us to treat clients similarly, and to define response criteria to which each subject can relate in a similar manner. They allow unprecedented control over experimentally relevant stimuli in operant conditioning and differentiation procedures and provide a seemingly endless horizon against which to cast our sights for extensions and applications. However, like all operant behavior, the benefits accruing from the use of percentile procedures can be delivered only after a response has been emitted. I trust that, relative to your recent history, you may find their operation reinforcing.

## REFERENCES

Alleman, H. D., & Platt, J. R. (1973). Differential reinforcement of interresponse times with controlled probability of reinforcement per response. *Learning and Motivation, 4,* 40–73.

Arbuckle, J. L., & Lattal, K. A. (1992). Molecular contingencies in schedules of intermittent punishment. *Journal of the Experimental Analysis of Behavior, 58,* 361–375.

Davis, E. R., & Platt, J. R. (1983). Contiguity and contingency in the acquisition and maintenance of an operant. *Learning and Motivation, 14,* 487–512.

Epstein, R. (1983). Resurgence of previously reinforced behavior during extinction. *Behaviour Analysis Letters, 3,* 391–397.

Galbicka, G. (1988). Differentiating *The Behavior of Organisms. Journal of the Experimental Analysis of Behavior, 50,* 343–354.

Galbicka, G., Fowler, K., & Ritch, Z. (1991). Control over response number by a targeted percentile schedule: Reinforcement loss and the acute effects of *d*-amphetamine. *Journal of the Experimental Analysis of Behavior, 56,* 205–215.

Galbicka, G., Kautz, M., & Jagers, T. (1993). Response acquisition under targeted percentile schedules: A continuing quandary for molar models of operant behavior. *Journal of the Experimental Analysis of Behavior, 60,* 171–184.

Galbicka, G., Kautz, M., & Ritch, Z. (1992). Reinforcement loss and behavioral tolerance to *d*-amphetamine: Using percentile schedules to control reinforcement density. *Behavioural Pharmacology, 3,* 535–544.

Galbicka, G., & Platt, J. R. (1986). Parametric manipulation of interresponse-time contingency independent of reinforcement rate. *Journal of Experimental Psychology: Animal Behavior Processes, 12,* 371–380.

Galbicka, G., & Platt, J. R. (1989). Response-reinforcer contingency and spatially defined operants: Testing an invariance property of phi. *Journal of the Experimental Analysis of Behavior, 51,* 145–162.

Holland, J. G., & Skinner, B. F. (1961). *The analysis of behavior.* New York: McGraw-Hill.

Kuch, D. O., & Platt, J. R. (1976). Reinforcement rate and interresponse time differentiation. *Journal of the Experimental Analysis of Behavior, 26,* 471–486.

Machado, A. (1989). Operant conditioning of behavioral variability using a percentile reinforcement schedule. *Journal of the Experimental Analysis of Behavior, 52,* 155–166.

McDade, C. E., & Goggans, A. (1993). Computer-based precision learning: Achieving fluency with college students. *Education and Treatment of Children, 16,* 290–305.

Platt, J. R. (1973). Percentile reinforcement: Paradigms for experimental analysis of response shaping. In G. H. Bower (Ed.), *The psychology of learning and motivation: Vol. 7. Advances in theory and research* (pp. 271–296). New York: Academic Press.

Platt, J. R. (1984). Motivational and response factors in temporal differentiation. In J. Gibbon & L. Allan (Eds.), *Timing and time perception* (pp. 200–210). New York: New York Academy of Sciences.

Platt, J. R., Kuch, D. O., & Bitgood, S. C. (1973). Rats' lever-press durations as psychophysical judgments of time. *Journal of the Experimental Analysis of Behavior, 19,* 239–250.

Scott, G. K., & Platt, J. R. (1985). Model of response-reinforcer contingency. *Journal of Experimental Psychology: Animal Behavior Processes, 11,* 152–171.

Skinner, B. F. (1938). *The behavior of organisms.* New York: Appleton-Century-Crofts.

Skinner, B. F. (1953). *Science and human behavior.* New York: MacMillan.

*Received July 6, 1993*
*Initial editorial decision October 3, 1993*
*Revisions received January 4, 1994; March 2, 1994; April 22, 1994*
*Final acceptance May 10, 1994*
*Action Editor, F. Charles Mace*